EXHIBIT 28

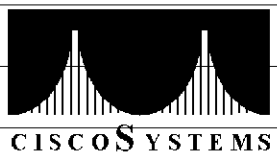

Document Number	ENG-120379
Revision	A
Authors	Massimiliano Ardica Saci Nambakkam
Project Manager	Sridhar Aswathnarayan

IOS 802.3ad on Constellation

Software Unit Functional Specification

Project Headline

This document describes the functionality provided by IEEE 802.3ad specification on Constellation platform.

Reviewers

Department	Name	Acceptance Date
Eng	Sridhar Aswathnarayan	

Modification History

Rev.	Date	Originator	Comment

Table of Contents

Problem Definition	3
Goals	3
Functional Description	3
Design Considerations	3
Software Interactions	6
Command Line Interface	8
Control Flow	16
SNMP	17
Differences with CatOS	17
Differences between Layer2 and Layer3 channels	17
Syslog	17
Timeframe	18
Testing Considerations	18

Definitions

This section defines words, acronyms, and actions which may not be readily understood.

AgPort	Aggregated port in the channel
Cosmos	IOS implementation on Constellation platform
CLI	Command Line Interface
FEC	Fast Ether Channel
FSM	Finite State Machine
LACP	Link Aggregation Control Protocol
PAgP	Port Aggregation Protocol
RPC	Remote Procedure Call
RP	Route Processor (also known as Draco)
SP	Switch Processor (also known as NMP)

1.0 Problem Definition

Ethernet channeling is the aggregation of multiple ports into one logical port to provide a higher aggregated bandwidth and link redundancy. The 802.3ad bundles ports with similar characteristics to form a channel through dynamic negotiation with the partners, ensuring their bidirectional and point-to-point connections. Cosmos will support a maximum of eight ports in each channel. The protocol is similar to PAgP [1].

The Marker protocol will also be implemented and will follow the same flow as that of LACP. High availability issues are not considered here.

2.0 Goals

Form channels with devices running 3ad (which include CatOS based constellation boxes, native constellation boxes as well as third party vendors)

Allow multiple channels to coexist simultaneously in the system.

Support of PAgP and 802.3ad in the same constellation system.

Support bundling of L3 ports.

Maintain the same interface/functionality with existing PAgP code and maximize code-reuse.

Consistent with the Cosmos software architecture; platform independent protocol, platform dependent layer and Port Manager and communication between the various subsystems defined.

3.0 Functional Description

Users specify the list of ports which are to be aggregated by specifying them to be in the same channel group. If all those ports share the same characteristics they are aggregated to provide a unified link capable of the sum of their bandwidth. A port can be a member of only one group.

If the channel mode is set to "on", the channel is bundled without running the LACP protocol. If the mode is "active" (equivalent to "desirable" in PAgP), LACP initiates protocol packets on that interface and attempts to form the bundle with the rest of the specified ports. If the mode is "passive" (equivalent to "auto" in PAgP), the port will participate in formation of the bundle only when it receives LACP protocol packets from its partner system.

Layer 2 (MAC based) and Layer 3 (IP based) channels can be formed by LACP. See section 10 for differences.

LACP is not compatible with PAgP. Also a port which is specified in LACP cannot be specified in PAgP or vice versa.

4.0 Design Considerations

4.1 Memory

The extra memory required for LACP is not much. Memory for the logical interface IDB's (hardware and software) and the subblocks for LACP interfaces are allocated dynamically and is not an overhead when it is not used.

One virtual idb would be allocated for every LACP channel. Two subblocks (LACP and FEC) are allocated for each interface participating in LACP channel. The LACP subblock contains protocol related information such as timers, partner information, state values, packet counters etc. The FEC subblock will have group_capability of the port, admin_group, channel mode etc. The total memory for the IDB and subblocks is expected to be of the order of 500 bytes.

4.2 Processing

The processing involved for the LACP is not huge. The system performance will not be affected by this feature.

There could be a slight delay before a port can participate in STP. The delay is for the LACP protocol to go over its FSM on the port. If the port is aggregated, a single instance of STP is run over the channel, rather than on each individual port in the channel.

4.3 Processors

The LACP code primarily runs on the SP. The user interface as well as the SNMP resides on the RP. Program control flow between the processors takes place by means of RPC. It is transparent to processes running on the system. The information flow between the (processes on the) processors take place via IOS IPC.

4.4 Image Bundling

The LACP component is not a mandatory part of the IOS image. Hence it will be made independent of the other sub-systems, such that it is possible to make a build without the LACP functionality. To facilitate this, all interactions from any other components to the LACP code is done by means of using Registries. This would facilitate that the LACP API's would get invoked if it is present as part of the image else they would return as a successful execution of the API. The LACP commands will also not be available if it is not part of the image.

4.5 Platform dependent and independent code

The LACP will be split into the platform independent (LACP Process) and the dependent portions (LACP Layer). The platform dependent portion will utilize the PAgP Shim (which includes FEC) for performing most of its functions. The SNMP subsection (MIB implementation) would be part of the platform independent portion.

The LACP Layer isolates the LACP process from the Port Manager (PM). Port Manager informs the LACP layer when port/interface attributes change. It also deals with control flow from NMP to the line cards.

The LACP process contains the protocol's state machine. It is responsible for transmitting/receiving/processing LACP PDUs. It gets the necessary information from the LACP Layer and after making the necessary computations informs the LACP layer whether ports are to be bundled.

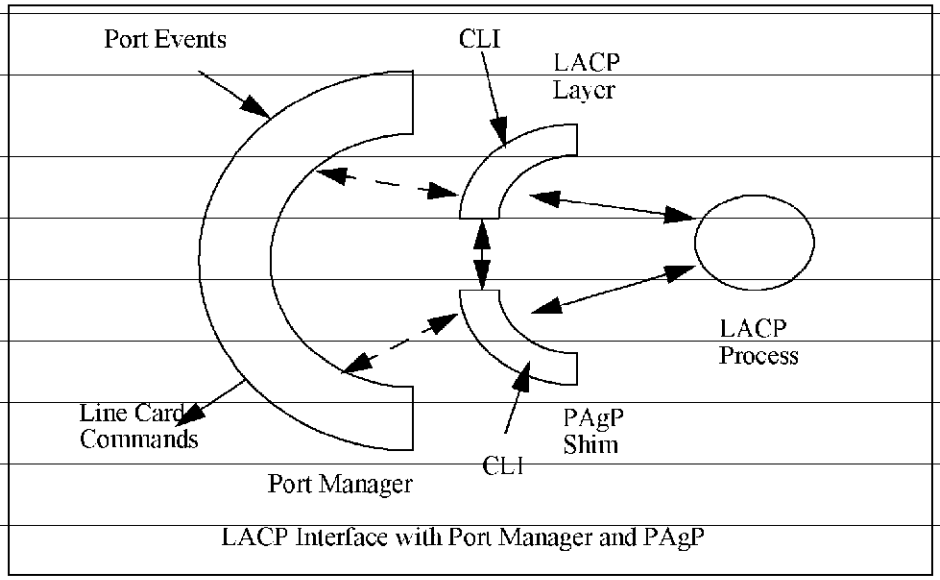

LACP Interface with Port Manager and PAgP

Port properties, namely the constraints define the capability of a port. Based on the Port's Speed, Duplex, Native VLAN, VLAN Ranges, Trunking Type and the admin-group in which it is, an operational Key is assigned to the port. This key is exchanged with its Partner's LACP process by means of exchanging LACP PDU's. If the capabilities of the ports match, and the mode is appropriate, the bundle can be formed.

The LACP layer does not need a process to handle its functions. It executes its functions from the calling context of the task, whereas the LACP Process is spawned initially and is in a loop waiting for events such as timer expiry, incoming packets to trigger its states. The PAgP shim currently utilizes a light weight process and the LACP layer can make use of the shim process to handle any specific tasks in future.

4.6 LACP Channel Modes

The LACP channel can be assigned one of the 4 possible modes:- On, Off, Active and Passive. They are analogous to the On, Off, Desirable and Auto of PAgP. In IOS if the port is not configured for a channel, the mode of the port is Off. Assuming the port capabilities being the same, the following combinations are possible

Table 1: Channel mode possibilities

Actor port mode	Partner port mode	Bundling Result (in Actor)
Off (channel not configured)	*	Bundle not formed
ON	*	Bundle formed
Active	Active/Passive	Bundle formed
Active	On/Off	Bundle not formed
Passive	Active	Bundle formed
Passive	On/Off/Passive	Bundle not formed
Active/Passive	Pagp Auto/ Desirable	Bundle not formed

4.7 Kernel resources

The Protocol would make use of system provided features such as timers to keep track of the state machine states. The LACP would utilize the Inband driver to send PDU's via the ports which LACP enabled. Line card communication is isolated by means of Port Manager.

5.0 Software Interactions

5.1 Interfaces with PAgP shim/LACP layer with the PM

The LACP Layer interacts with the PAgP shim to leverage the existing code. PAgP shim is the platform dependent portion of the PAgP protocol. LACP can exist without PAgP bundled in the image, However PAgP shim should be present in the image for the proper functioning of the LACP protocol.

Events such as port link state changes and port parameter changes are generated by the PM. Module inserted and removed are translated into multiple ports up/down events. The PM provides all the properties of the port which are used in the decision making process of deciding to bundle the port or not. PM will notify the LACP layer of any CLI/SNMP changes to the port properties (port speed, VLAN configuration etc.), or any link state changes. The control from PM to LACP layer would be in the form of registry calls.

LACP layer/PAgP shim also need to communicate its information to the PM (e.g. to perform a soft link down or up, so that shim can follow the same flow when it needs to bundle a port or remove it, configure port bundle selects, etc.).

The LACP relies on the following interface which exists between PAgP shim and PM

reg_invoke_pm_port_bndl_start(port_cookie pc)

This message from PM to PAgP shim informs that the link on this interface has come up.

reg_invoke_pm_port_bndl_stop(port_cookie pc);

This message from PM to PAgP shim informs that the link on this interface is down. This generates the link down processing, however the interface is not removed from the channel.

reg_invoke_pm_port_statechange(pm_port_reason_t pr,
pm_port_state_t ps, port_cookie *pc, int cookie_count)

This message from PM indicates that port state has changed to present from not present (or vice versa) resulting in an equivalent port up/down effect. This could happen when the user changed the port between layer2 and layer3.

reg_invoke_ec_port_attrib_change(idbtype *swidb)

This message from PM indicates that port attributes have changed for this interface. This could cause new port to be bundled because of the attribute change.

reg_invoke_ec_port_attrib_diff(idbtype *swidb1, idbtype *swidb2, boolean admin,
boolean oper, boolean inform, char *reason)

This API compares the etherchannel related port attributes like channel group and mode (auto/desirable). In the case of LACP, the new modes will be included.

reg_invoke_hwif_comingup(hwidbtype *idb)

reg_invoke_hwif_goingdown(hwidbtype *idb)

These two API's are invoked when the logical interface changes from/to shut/noshut. This is usually followed by the port up/down messages for the individual interfaces.

The group capability (GC) related functionality of PAgP shim is not utilized. Other registry functions such as checking if LACP is available will be provided as needed. New functionality such as Standby ports in LACP will be part of the LACP layer.

5.2 Interface of LACP layer with the LACP process

When the LACP layer decides that the interface matches all the requirements and it can be bundled, it informs the LACP process to initiate the protocol. The API calls are by means of registries since the image could be made to support channel mode "on" but not the protocol itself (active/passive) by not including it in the image. The registry calls from the LACP layer to the LACP process are as follows:-

reg_invoke_lacp_disable (idbtype *swidb, fec_mode_t)

This API will be called from the LACP layer when the interface is removed from the channel. The protocol stops running on this interface.

reg_invoke_lacp_enable (idbtype *swidb)

This API will be called from the LACP layer when the interface is added to the channel. The protocol is started on this interface.

reg_invoke_lacp_portdown (idbtype *swidb)

This API will be called when the LACP layer decides that the port cannot be bundled due to a change in the port characteristics. The port will be removed from the Aggregator.

reg_invoke_lacp_portup (idbtype *swidb)

This API will be called when the LACP layer concludes that the port may be part of the Aggregator.

5.3 Interfaces between PAgP Process and the PM

The PM will not invoke any API's of the LACP process directly. The LACP makes use of the generic (non-registry based) PM API's defined in the file pm_port_public.h. Some of those are pm_port_bndl_enable(), pm_port_get_swidb(), pm_port_get_cookie_from_swidb(), pm_port_l3_agport_list() and pm_port_remove_from_bundle().

5.4 Other Interfaces

The other protocols interact with LACP in the same way as they do with PAgP, namely:-

The logical interface will be transparent for the protocols like CDP, DTP to transmit/receive packets over physical interfaces though the interfaces belong to a bundle. On the trunk ports, LACP PDUs will be transmitted/received on the lowest numbered VLAN. STP, VTP will have to deal with the logical-port only. STP packets are sent over the first port of the bundle.

The MAC address of the L3 bundle will be the MAC address of the first port of the admin-group, irrespective of the port being down.

LACP will transmit/receive LACP PDUs only from interfaces that are 'up' and have LACP enabled. The following API is called by the Inband driver to enque packets to the LACP queue, so that the LACP process can process the LACP PDU's.

reg_invoke_raw_enqueue (LINK_LACP, lacp_enqueue, "lacp_enqueue");

6.0 Command Line Interface

The CLI available pertaining to channelling can be distinguished in two sets:

- commands to configure and display generic informations for the channels present in the system and
- commands specific for the protocol (PAgP or LACP) enabled on these channels

The first set of commands are already implemented and they will be only slightly modified to accomplish the introduction of the new LACP protocol (the additional information will be highlighted in bold); the commands belonging to the second set will be introduced for the first time, having as a reference the PAgP's specific CLI (see IOS User Guide).

6.1 Generic commands

6.1.1 Command to add an interface to a channel

```
Router(config-if)#channel-group ?
    <1-256>  Channel group number
```

The total number of channel groups allowed in the system is **256**. The channel group number is global and is shared between all the channeling protocols. If a specific channel number is used for a channel group whose interfaces have enabled PAgP, that same channel number can not be used for configuring a channel having interfaces have LACP enabled or vice versa.

```
Router(config-if)#channel-group 1 mode ?
    auto        Enable PAgP only if a PAgP device is detected
    desirable   Enable PAgP unconditionally
    passive     Enable LACP only if a LACP device is detected
    active      Enable LACP unconditionally
    on          Enable Etherchannel only
```

The first two options (auto and desirable) are valid only for interfaces on which PAgP protocol is enabled and the command would be rejected if issued on interfaces on which LACP is enabled.

In the same way, the third and fourth options (active and passive) are valid only for interfaces on which LACP is enabled.

The last option "on" will force the bundle of the interface on the channel without any kind of negotiation (neither PAgP or LACP process will be spawned for that interface).

NOTE: in this implementation, the user does not need to specify the Administrative Key to be assigned to the specific interface (different from the CatOS approach [2]); the value of the Administrative Key will instead be picked up automatically by the system, if LACP is enabled on that interface. Of course all the interfaces added to the same channel-group (i.e. having the same channel number specified in the command above) will be assigned the same Administrative Key value. This means that the system will pick up a value, when the first interface is added to the channel (i.e. when a new channel group id is specified), and this value will be propagated to all the interfaces added later to the same channel.

6.1.2 Command to select the load-balance method

```
Router(config)#port-channel load-balance ?
    dst-ip          Dst IP Addr
    dst-mac         Dst Mac Addr
    src-dst-ip      Src XOR Dst IP Addr
    src-dst-mac     Src XOR Dst Mac Addr
```

```
src-ip          Src IP Addr
src-mac         Src Mac Addr
```

This global configuration command is used to select the channel load balancing method. The methods can be MAC (layer 2) or IP (layer 3). Furthermore, it can based on the source address, destination address or both source and destination addresses. Load-balancing cannot be changed on a per channel basis.

6.1.3 Command to display brief etherchannel info

```
Router>sh etherchannel brief
                Channel-group listing:
                ----------------------

Group: 1
----------
Group state = L2
Ports: 4   Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol: LACP
```

The last line was added for each channel to specify the protocol enabled for the physical members of the channel group. If the channel has been created in ON mode, a (-) is displayed in that field.

6.1.4 Command to display detailed etherchannel info

```
Router>sh etherchannel detail
                Channel-group listing:
                ----------------------

Group: 1
----------
Group state = L2
Ports: 4   Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol: LACP
                Ports in the group:
                -------------------
Port: Fa4/1
------------

Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Active        Gcchange = -
Port-channel  = Po1         GC   =   -           Pseudo port-channel = Po1
Port index    = 0           Load = 0x55          Protocol = LACP

Flags:  S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs
        A - Device is in active mode       P - Device is in passive mode

Local information:
                          LACP Port     Oper     Port
Port      Flags State     Priority      Key      State
Fa4/1      SA     bndl       128         100      0x75

Partner's information:

          Partner           Partner           Partner        Partner
Port      Name              System ID         Port     Age   Flags
Fa4/1     TBA04010428    8000,00b0.c23e.d861   3/1     14s    SP

          Partner       Partner
          Oper Key     Port State
            128          0x81

Age of the port in the current state: 16h:27m:42s
Port: Fa4/2
------------

Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Active        Gcchange = -
Port-channel  = Po1         GC   =   -           Pseudo port-channel = Po1
Port index    = 1           Load = 0xAA          Protocol = LACP

Flags:  S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs
```

```
        A - Device is in active mode       P - Device is in passive mode

Local information:
                            LACP Port      Oper     Port
Port      Flags State       Priority       Key      State
Fa4/1      SA    bndl         128          100      0x75

Partner's information:

          Partner           Partner           Partner        Partner
Port      Name              System ID         Port     Age   Flags
Fa4/2     TBA04010428       8000,00b0.e23e.d861  3/2   20s    SP

          Partner       Partner
          Oper Key      Port State
            128           0x81

Age of the port in the current state: 16h:28m:55s

Port-channel: Po1
------------

Age of the Port-channel   = 143h:01m:12s
Logical slot/port   = 14/1            Number of ports = 2
GC                  = -               HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            = LACP

Ports in the Port-channel:

Index   Load   Port     EC state
------+------+------+------------
  0     55     Fa4/1    active
  1     AA     Fa4/2    active

Time since last port bundled:    16h:28m:58s    Fa4/1
Time since last port Un-bundled: 16h:29m:00s    Fa4/4
```

6.1.5 Command to display the load-balance method

```
Router#sh etherchannel load-balance
Source XOR Destination IP address
```

This command shows the load balanced logic configured in the system; the possible alternatives are described in paragraph 6.1.2.

6.1.6 Command to display port info

```
Router>sh etherchannel port
                Channel group listing:
                ----------------------

Group: 1
----------
                Ports in the group:
                -------------------
Port: Fa4/1
------------

Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Active         Gcchange = -
Port-channel  = Po1         GC   =   -            Pseudo port-channel = Po1
Port index    = 0           Load = 0x55           Protocol = LACP

Flags:  S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs
        A - Device is in active mode       P - Device is in passive mode

Local information:
                            LACP Port      Oper     Port
Port      Flags State       Priority       Key      State
Fa4/1      SA    bndl         128          100      0x75

Partner's information:
```

```
          Partner           Partner           Partner        Partner
Port      Name              System ID         Port      Age  Flags
Fa4/1     TBA04010428    8000,00b0.c23e.d861   3/1      14s   SP

          Partner       Partner
          Oper Key     Port State
            128          0x81

Age of the port in the current state: 16h:27m:42s
Port: Fa4/2
------------

Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Active        Gcchange = -
Port-channel  = Po1         GC   = -             Pseudo port-channel = Po1
Port index    = 1           Load = 0xAA          Protocol = LACP

Flags:  S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs
        A - Device is in active mode       P - Device is in passive mode

Local information:
                          LACP Port     Oper     Port
Port      Flags State     Priority      Key      State
Fa4/1      SA     bndl       128        100      0x75

Partner's information:

          Partner           Partner           Partner        Partner
Port      Name              System ID         Port      Age  Flags
Fa4/2     TBA04010428    8000,00b0.c23e.d861   3/2      20s   SP

          Partner       Partner
          Oper Key     Port State
            128          0x81

Age of the port in the current state: 16h:28m:55s
```

The output of this command is very similar to the one described in 6.1.4, with the difference being that only information related to the ports are displayed.

6.1.7 Command to display port-channel info

```
Router>show etherchannel port-channel
                Channel-group listing:
                ----------------------

Group: 1
----------
                Port-channels in the group:
                ---------------------

Port-channel: Po1
------------

Age of the Port-channel   = 143h:01m:12s
Logical slot/port   = 14/1           Number of ports = 2
GC                  = -              HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            = LACP

Ports in the Port-channel:

Index   Load   Port     EC state
------+------+------+------------
  0      55     Fa4/1     active
  1      AA     Fa4/2     active

Time since last port bundled:    16h:28m:58s    Fa4/1
Time since last port Un-bundled: 16h:29m:00s    Fa4/4
```

The info displayed with this command are similar to that in 10.1.4, with the only difference that only informations related to the port-channels will be described.

6.1.8 Command to display summary info

```
Router>sh etherchannel summary
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        R - Layer3      S - Layer2
Group Port-channel  Protocol  Ports
------+-------------+---------+-----------------------------------------------
1     Po1(SU)         LACP      Fa4/1(P)    Fa4/2(I)    Fa4/3(P)    Fa4/4(s)
2     Po2(RU)         PAGP      Fa3/1(P)    Fa3/3(P)
10    Po10(SU)          -       Fa6/2(P)
```

The column to display the protocol used for the channel has been added to the command; if the channel mode is ON, a (-) will be displayed.

6.1.9 Command to show info for the logical interface

```
Router>show interface port-channel 1 protocol
LACP
```

This command displays the protocol used on the specified portchannel; this depends on the protocol enabled on the interfaces members of the channel.

Note: if the interface are configured as part of the channel in mode ON, the command would display

```
Protocol: -                 Mode: ON
```

```
Router>sh int port-channel 1 etherchannel
Age of the Port-channel   = 126h:00m:15s
Logical slot/port   = 14/1          Number of ports = 2
GC                  = 0x00000000    HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol running    = LACP

Ports in the Port-channel:

Index   Load   Port    EC state
------+------+------+------------
  0     55     Fa4/3   active
  1     AA     Fa4/4   active

Time since last port bundled:    00h:06m:35s    Fa4/4
Time since last port Un-bundled: 00h:27m:10s    Fa4/4
```

The informations displayed with this command is the same as that of 6.1.7 (show etherchannel port-channel)

6.1.10 Command to show the info for a physical member

```
Router>show int fa4/1 etherchannel
Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Active        Gcchange = -
Port-channel  = Po1         GC   =  -            Pseudo port-channel = Po1
Port index    = 0           Load = 0x55          Protocol = LACP

Flags:  S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs
        A - Device is in active mode       P - Device is in passive mode

Local information:
                            LACP Port     Oper     Port
Port      Flags State       Priority      Key      State
Fa4/1      SA     bndl        128         100      0x75

Partner's information:

          Partner           Partner           Partner        Partner
Port      Name              System ID         Port     Age   Flags
Fa4/1     TBA04010428     8000,00b0.c23e.d861   3/1     14s    SP

          Partner       Partner
          Oper Key      Port State
```

```
                128             0x81

Age of the port in the current state: 16h:27m:42s
```

The output of this command is the same as 6.1.6 but concerning only the specified interface.

6.2 LACP specific CLI commands

6.2.1 Command to set the system priority

```
Router(config)#lacp system-priority ?
  <1-65535>   System priority
```

This command sets the priority of the system. The values allowed are in the range [1..65535]: the higher the number, the lower the priority. The default priority is **32768**.

Although this is a global configuration command, the parameter will be used only for configuring port-channel having physical interfaces with LACP protocol enabled, whereas it will be ignored in the other cases.

6.2.2 Command to choose the protocol

```
Router(config-if)#channel-protocol ?
  pagp   Enable PAgP protocol for the interface
  lacp   Enable LACP protocol for the interface
```

This command sets the protocol used on the specific interface to manage channelling. The protocol used by default is **PAgP**. The command will be rejected if issued on an interface which is already belonging to a channel. Hence if the user wishes to change the channel protocol for the interface, then the interface should be removed from the channel group and then added to the new channel group and the new channel protocol specified.

6.2.3 Command to set the port priority

```
Router(config-if)#lacp port-priority ?
  <1-255>   Port priority
```

This command sets the priority for the physical interfaces. The values allowed are in the range [1..255], where higher numbers correspond to lower priorities. The priority used by default is **128**.

This command can be issued only for interfaces on which LACP protocol is enabled, while it would be hidden in the other cases.

6.2.4 Command to show LACP System ID

```
Router>show lacp sys-id
8000,AC-12-34-56-78-90
```

This command is used to show the System Identifier adopted by LACP protocol. As the System ID is made up by the concatenation of the System Priority and the System MAC address this command can be used to know one of these parameters as well. The first two byte are the system priority, whereas the last six bytes are the globally administered individual MAC address associated to the system.

6.2.5 Command to show LACP statistics

```
Router#sh lacp 1 counters
          LACPDUs         Marker         LACPDUs
```

```
Port        Sent    Recv      Sent    Recv      Pkts Err
-----------------------------------------------------
Channel group: 1
  Fa4/1     8       15        0       0         0
  Fa4/2     14      18        0       0         0
  Fa4/3     14      18        0       0         0
  Fa4/4     13      18        0       0         0
```

This command shows statistics for all the physical interfaces belonging to the same channel-group specified on the command line (or for all the port-channel configured in the system if no value is specified). The second and third column display the LACPDUs sent and received on each specific interface, whereas the last two columns show the marker protocol packets.

6.2.6 Command to show LACP internal info

```
Router#sh lacp 1 internal
Flags:  S - Device sends PDUs at slow rate. F - Device sends PDUs at fast rate.
        A - Device is in Active mode.       P - Device is in Passive mode.

Channel group 1
                            LACPDUs     LACP Port     Oper     Port
Port      Flags     State   Interval    Priority      Key      State
Fa4/1     saC       bndl      30s         128         100      0x75
Fa4/2     saC       bndl      30s         128         100      0x75
Fa4/3     saC       bndl      30s         128         100      0x75
Fa4/4     saC       bndl      30s         128         100      0x75
```

This command is used to display internal information for the interfaces belonging to the specified channel (or for all the port-channel configured in the system if no value is specified). In the **State** column, the state of the specific port at the current moment is displayed. Allowed values are:-

bndl -> the port is attached to an aggregator and bundled with other ports

susp -> the port is in suspended state, so is not attached to any aggregator

hot-sby -> the port is in Hot-standby state

down -> the port is down

The **Port State** represents the Actor's state variables for the port, encoded as individual bits within a single octet and having the following meaning [1]:

bit0: *LACP_Activity*

bit1: *LACP_Timeout*

bit2: *Aggregation*

bit3: *Synchronization*

bit4: *Collecting*

bit5: *Distributing*

bit6: *Defaulted*

bit7: *Expired*

6.2.7 Command to display LACP neighbors info

```
Router#sh lacp 1 neighbor
Flags:  S - Device sends PDUs at slow rate. F - Device sends PDUs at fast rate.
        A - Device is in Active mode.       P - Device is in Passive mode.

Channel group 1 neighbors
          Partner              Partner          Partner                    Oper    Port
Port      Name                 System ID         Port      Age   Flags     Key    State
Fa4/1     TBA04010428     8000,00b0.c23e.d84e    3/1       29s     P       200     0x81
Fa4/2     TBA04010428     8000,00b0.c23e.d84e    3/2        0s     P       200     0x81
Fa4/3     TBA04010428     8000,00b0.c23e.d84e    3/3        0s     P       200     0x81
```

```
Fa4/4      TBA04010428     8000,00b0.c23e.d84e   3/4        0s     P      200     0x81
```

This command is used to display information about the partner system for a specified port channel (or for all the port-channel configured in the system if no value is specified).

Note: the info displayed reflect the operational informations that the Actor has about the Partner; these are retrieved from the received LACPDUs. In the case where no PDUs have been received the default administrative information would be displayed in braces.

6.2.8 Debug CLI command

These commands are very useful for debugging and understanding the behavior of LACP.

```
Router#sh lacp 1 sm state
Channel group 1:

Port          RX SM         PERIODIC TX SM        MUX SM
Fa4/1         curr               fa_p               dis
Fa4/2         curr               fa_p               dis
Fa4/3         curr               fa_p               dis
Fa4/4         curr               fa_p               dis
```

The output displays the state in the different state machines for each physical interfaces belonging to the channel. The possible states are the following:

- RX State Machine

init = INITIALIZE
p_dis = PORT_DISABLED
exp = EXPIRED
l_dis = LACP_DISABLED
def = DEFAULTED
curr = CURRENT

- PERIODIC TX State Machine

no_p = NO_PERIODIC
fa_p = FAST_PERIODIC
sl_p = SLOW_PERIODIC
p_tx = PERIODIC_TX

- MUX State Machine

det = DETACHED
wait = WAITING
att = ATTACHED
coll = COLLECTING
dis = DISTRIBUTING

6.3 Enable/disable debug messages

```
Router#debug lacp ?
  all       All debugging
  detail    Step below all
  error     Displaying error messages
  event     Major events
  <cr>
```

7.0 Control Flow

The following figure illustrates the flow in the case of an interface added to channel and link coming up. The steps are illustrated by means of the numbers 1..5. Since the PAgP shim and LACP layer are tightly coupled, the calls between them are not shown. Calls from PM to PAgP shim and LACP Layer to LACP Process are registry based whereas calls from LACP Process to LACP layer and LACP Layer to PM are invoked directly.

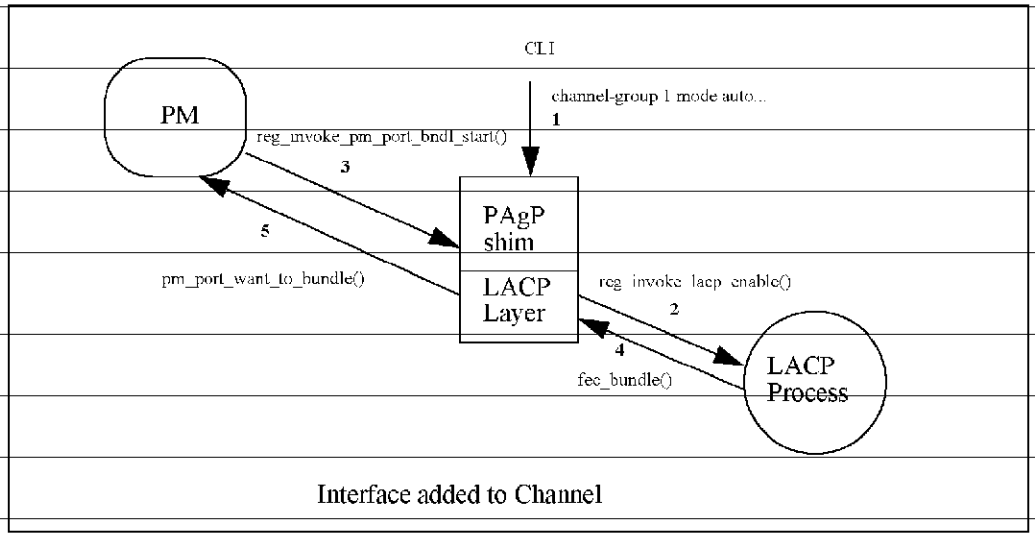

Interface added to Channel

If the port has just come up, the sequence would start from 2 and proceed to 5. At every step, it is checked if PAgP is enabled on that interface before proceeding and to avoid re-processing if interface is already up. Step 1 illustrates the user input. The LACP layer then allocates the FEC subblock, and updates it with the port parameters. On invocation of step 2, the LACP protocol is started and memory for the protocol specific subblock is allocated. Step 3 illustrates the link coming up. When the protocol decides that the port can be bundled, it invokes step 4, which updates the required structures and informs PM (step 5).

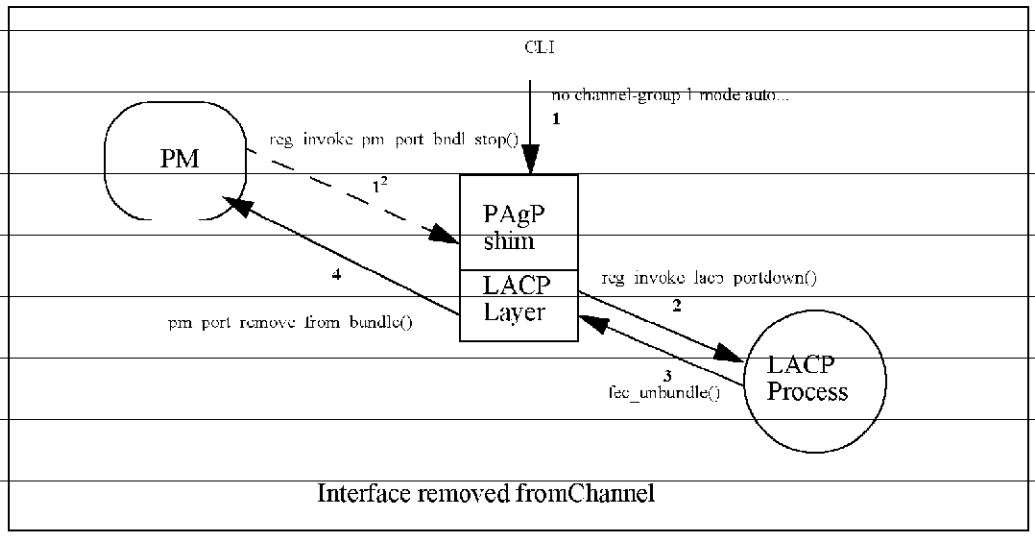

Interface removed fromChannel

When the user removes a port from the channel, the sequence of steps 2, 3 and 4 takes place. If the port is already down then steps 2 and 3 will not take place. The sequence 1^2 is invoked if the link goes down. Steps 2, 3 and 4 remain the same.

8.0 SNMP

The 802.3ad MIB as defined in [1] will be supported.

9.0 Differences with CatOS

The behavior of LACP is different from that of CatOS in the following cases:-

- The CatOS image provides support for forming multiple channels within the same channel group. The IOS LACP behavior is consistent with its PAgP counterpart and hence does not provide this support.
- When a change is affected on a port belonging to the channel in CatOS, the same change is applied on all the ports constituting the channel, whereas in the case of IOS, the dis-similarity causes the port to drop out of the Aggregator.
- The user does not need to select the Administrative Key to assign to the specific interface whereas it needs to be mentioned in CatOS.

10.0 Differences between Layer2 and Layer3 channels

By default the interface is layer3 and hence the channel formed is layer 3, unless explicitly specified as layer 2. There are a few subtle differences between the layer2 and layer3 channels in behavior.

- The MAC address assigned to the channel is that of the first port in the channel in the case of layer 3 whereas if the channel is layer 2 then the MAC address is that of the first Aggregator.
- If the channel is layer 2 and the partner is not running LACP protocol, then the port are treated as a normal switching port (I state). But if it is a layer 3 port, then the port are put in suspended state.

11.0 Syslog

The PAgP syslog messages will remain the same as they pertain to the etherchannel. The new ones specific to the LACP are:-

msgdef(LACP_BADQUEUEMSG, LACP, LOG_DEBUG, 0, "Unknown message in queue.");

This entry will be shown when a non-lacp type of message is received on the interface.

msgdef(LACP_PROCFAIL, LACP, LOG_ERR, 0, " Failed to initialise LACP process.");

This entry will be shown if the creation of LACP task failed during initialization time.

msgdef(LACP_NOMEM, LACP, LOG_ERR, 0, "No memory for PAgP process.");

This entry will be shown when there is no memory for LACP processes to allocate subblocks.

msgdef(LACP_QFULL, LACP, LOG_ERR, 0, "LACP queue is full.");

This entry will be shown when a queue used by LACP is full.

msgdef(LACP_INVALID_EVENT, LACP, LOG_ERR, 0, "Interface %s Current state %d Event type %d");

This entry will be shown when LACP Process encounters an invalid state or receives an invalid event. The current state and the event which occurred is displayed along with the interface name.

12.0 Timeframe

This feature will be available for the IOS release 12.1(7)E ?? Herschel?

13.0 Testing Considerations

Testing equipment should include a CatOS based Constellation 6500, and non-cisco equipment supporting 802.3ad

Reference Documents

1. IEEE Std 802.3ad-2000

2. Constellation Ethernet Channel with LACP: Software Unit Design ENG-97695

3. IOS PAgP/FEC on Constellation:Software Unit Functional Specification: ENG-26148